

NASA HPC Update

Supercomputers, Clusters & Hypercomputers (Past, Present Future)

by

Dr. Olaf Storaasli

Analytical & Computational Methods Branch


NASA Langley Research Center

Hampton Virginia

HPC Users Forum

Dearborn, Michigan

April 12-14 , 2004





Contents

Past HPC at NASA & Langley

Research: Langley, Ames, Glenn

Ops: JPL, Goddard, Dryden, Johnson, Kennedy, Marshall, MTC

Present Production (Cosmo + Center Clusters)

Research (CiCT + Center Clusters)

Future Clusters

iNASA


Hypercomputers....

GPS Factor/Solve Time: $[K]\{u\}=\{f\}$


SGI \Rightarrow **G5**




54,870
Equations




263,574 DOF



88,404 DOF
172,400 DOF




111,893 DOF



217,918 DOF

DOF/Eqns.	Terms $\neq 0$	B-OPS	SGI/0.25G	G4/1GHz*	G5/2GHz	MF	Speedup
54,870	1.3 M	3.8	25.9 / 1.0	14.3 / 0.5	7.18 / 0.13	529	3.6 / 7.7
88,404	1.7 M	7.7	49.8 / 1.0	24.5 / 0.7	8.86 / 0.19	866	5.6 / 5.2
111,893	1.6 M	14.6	236.8 / 2.7	125.7 / 1.4	40.8 / 0.33	358	5.8 / 8.2
172,400	7.2 M	12.7	133.7 / 3.3	60.2 / 2.2	21.7 / 0.37	585	6.2 / 8.9
217,918	5.7 M	41.2	230.7 / 8.3	116.2 / 3.2	54.27/1.09	758	4.3 / 7.6
263,574	6.3 M	17.6	166.4 / 3.9	78.1 / 2.6	26.1/0.134	675	6.4 / 29



Error: $\|Ku - f\| < 10^{-10}$

Langley Research Center

*Laptop = Cray C90

3

Olaf.O.Storaasli@nasa.gov HPCUF 4-04

Contents

Past HPC at NASA & Langley

Research: Langley, Ames, Glenn (Research)

Ops: JPL, Goddard, Dryden, Johnson, Kennedy, Marshall, MTC

Present Production (Cosmo + Center Clusters)

Research (CiCT + Center Clusters)

Future Clusters

iNASA

Hypercomputers....


A Revolution is Brewing in HPC

Accelerating Scientific Discovery and Understanding

The NASA Advanced Supercomputing (NAS) Division, working with NASA's Aeronautics and Earth Sciences Enterprises, has begun exploring new ways of doing high performance computing. The effort is focused on providing adequate HPC platforms with high availability to mission critical problems, in order to significantly accelerate the science discovery process in those areas.

Earth Sciences - The ECCO Project

The Consortium for Estimating the Circulation and Climate of the Ocean (ECCO) is a joint venture between Jet Propulsion Laboratory, MIT, and the Scripps Institute of Oceanography. A major current effort is to execute a number of decadal ocean simulations using MIT's MITgcm code at 1/4 degree global resolution, or better. This effort is the first major project to use the new SGI Altix 512p single-system image (SSI) computer at NAS. NASA's new HPC system has provided a significant increase in throughput to the ECCO team. Some hint of this is seen in the performance chart (left). Additional optimizations are planned for this code on the Altix platform. These efforts are expected to increase performance to almost a simulated decade per day on 512 CPUs. This level of computing capability, dedicated to a core science team, can revolutionize the rate of scientific discovery, a process critical for national leadership in the sciences.




SGI Altix - 512p SSI System

The SGI Altix SSI system has been operational since 10/30/03. It is routinely scaling production applications to 512 CPUs with excellent results.



MITgcm Performance - 1/4 Degree
(yrs/day vs CPU Count)




Aerosciences - The OSP and RTF Projects

A second major NAS effort to accelerate the time to solution for key mission programs, centers around work in the Aerosciences. NASA, in particular, needs extensive computational resources to address the specific needs of its Return to Flight (RTF) and Orbital Space Plane (OSP) projects. The NAS Division is working with NASA Enterprise leadership to define a small series of focused efforts in support of these activities. In the past, the OVERFLOW code has been essential for NASA and Aero industry needs in flow simulations over existing and notional designs. The Altix system has significantly improved the time to solution for this code. The chart (right) shows a comparison with previous best efforts.




Cray X1 - 16p Supercomputer


The NAS team will be using this 192 GFLOP/s system to explore the performance envelop of "new age" vector supercomputers on a variety of NASA codes.



OVERFLOW - 35M Point Problem
(GFLOP/s vs CPU Count)



+ VT (Xserve) & Hypercomputers




Langley Research Center


New Solution Methods

- Cellular Automata: Stephen Wolfram - *A New Kind of Science*
- Complexity via simple interactions w/o PDEs
- CFD => Structures



- Cell interacts with neighbors; one processor per cell



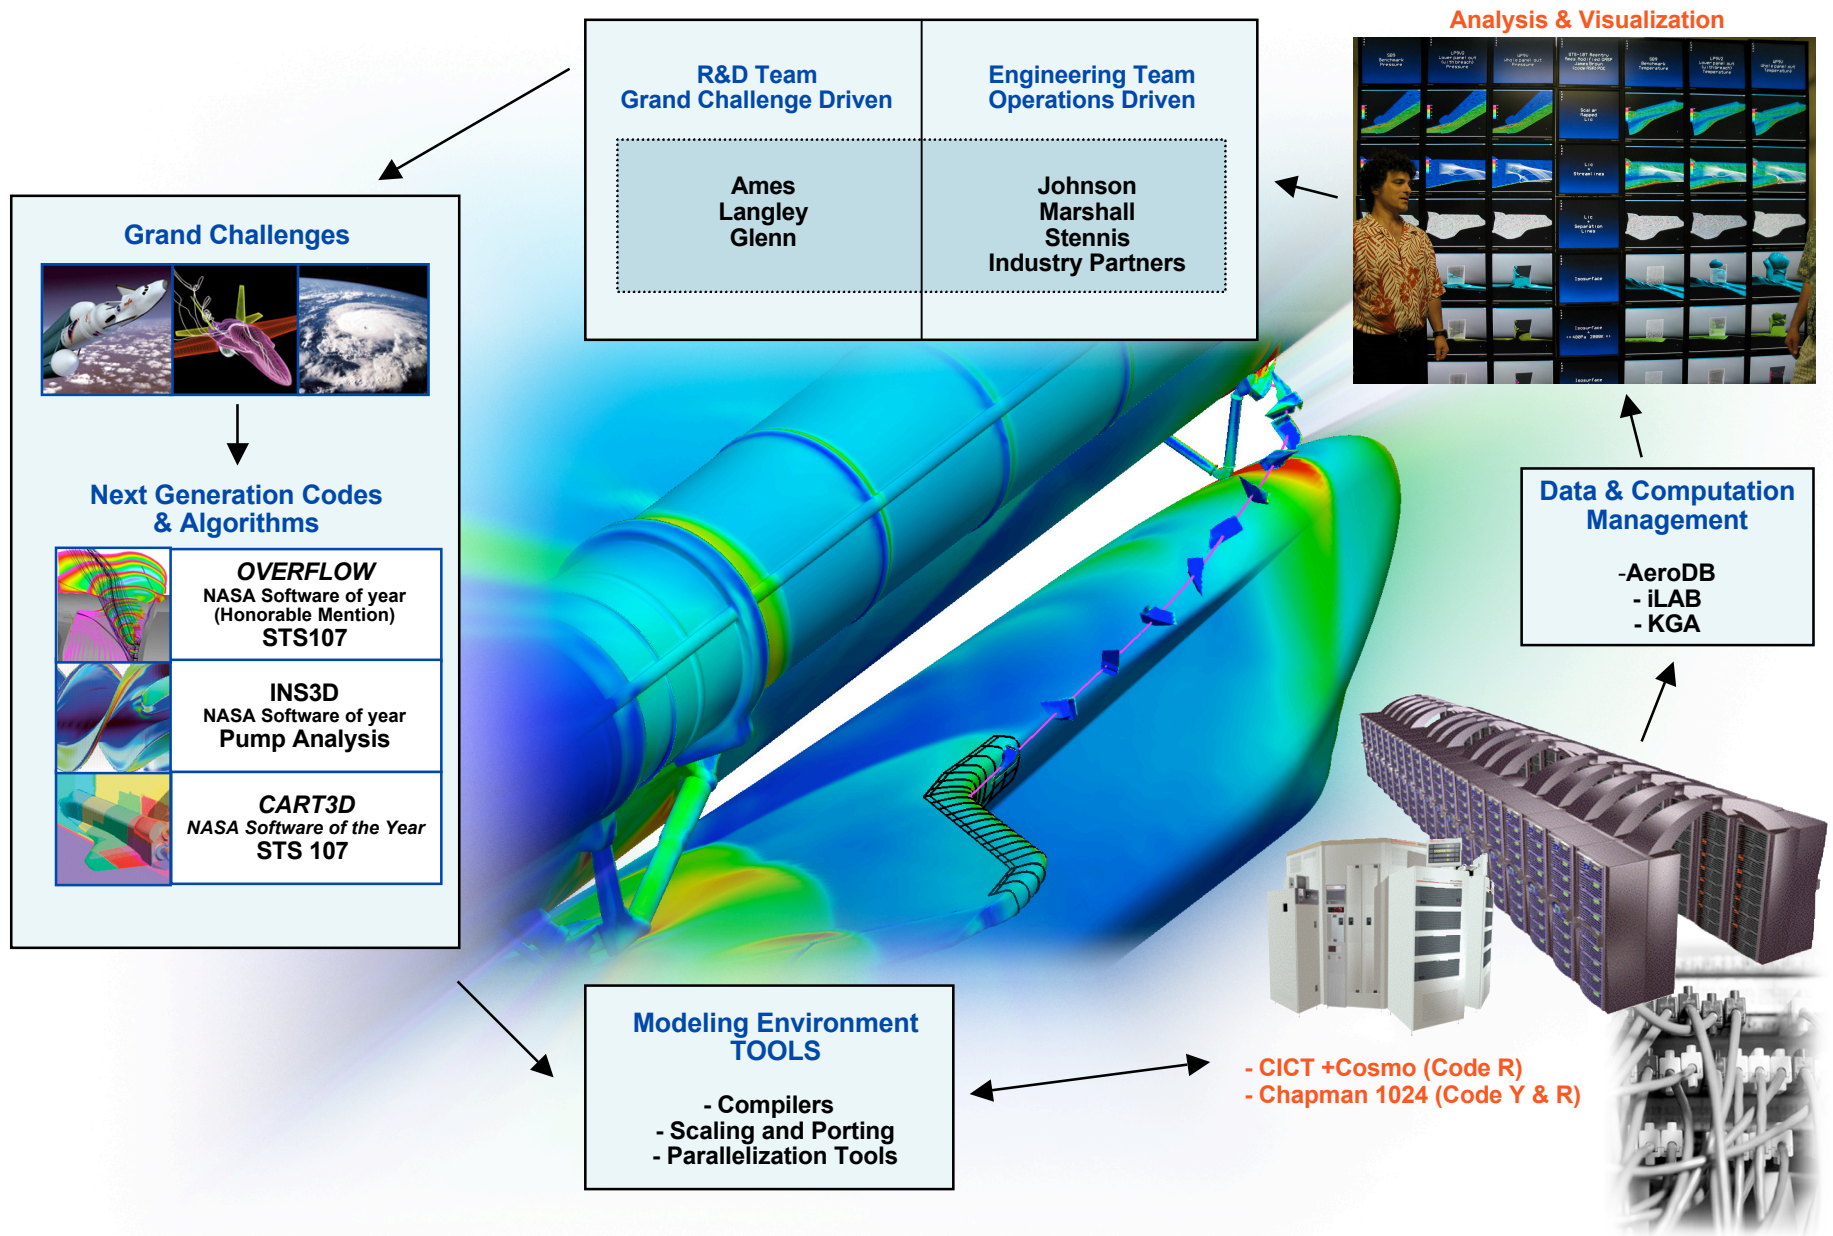
FEA
solution



Cellular
Automata
solution



STS-107 Investigation Support



NASA-wide HPC Systems




***SGI Altix: 128 Processors
(IA64 - 6 GFLOP/CPU)***



***Cray X1 : 64 Processors
(Custom - 12 GFLOP/CPU)***

+ 128 processor Opteron System





Contents

Past NASA & Langley HPC

Research: Langley, Ames, Glenn (Research)

Ops: JPL, Goddard, Dryden, Johnson, Kennedy, Marshall, MTC

Present Production (Cosmo + Center Clusters)

Research (CiCT + Center Clusters)

Future Clusters

iNASA

Hypercomputers....

Langley Cluster Growth


*Estimate by Joseph.H.Morrison@nasa.gov
Chair of Langley Cluster SIG*

1200 - 1500 nodes on 2/04
440 nodes added since 6/03
50% increase in nodes/year
for several years

Other NASA Centers: Similar growth?

Why? 3x – 20x less to compute locally
Cost/CPU hour roughly 1/10

Chair, Langley Cluster SIG



Growing HPC Interest at Langley



System X: Building the Virginia Tech Supercomputer

by Prof. Srinidhi Varadarajan

February 3, 2004, 2:00 P.M. in the H.J.E. Reid Auditorium.



Rank	Site Country / Year	Computer / Processors Manufacturer	R_{max} R_{peak}
1	<u>Earth Simulator Center</u> Japan/2002	Earth-Simulator / 5120 NEC	35860 40960
2	<u>Los Alamos National Laboratory</u> United States/2002	ASCI Q - AlphaServer SC45, 1.25 GHz / 8192 HP	13880 20480
3	<u>Virginia Tech</u> United States/2003	X 1100 Dual 2.0 GHz Apple G5/Mellanox Infiniband 4X/Cisco GigE / 2200 Self-made	10280 17600


**Many NASA G5 Xserve orders
Port of CFD & FE codes**



**The
Sigma
Series**

A Forum on Science & Technology

Sponsored by: NASA Langley Research Center, Virginia Air and Space Center, and Daily Press, Inc.



System X: *Building the Virginia Tech Supercomputer*

Dr. Srinidhi Varadarajan

**February 3, 2004
7:30 p.m.**

Virginia Air and Space Center

System X was conceived in March, designed in July and, by October 2003, it achieved a sustained performance of 10.28 Teraflops, thus becoming the World's third-fastest supercomputer. System X has several novel features. It is the first scientific supercomputer with Apple technology, using 1100 Apple G5 platforms which employ the new IBM PowerPC 970 64-bit CPU. Secondly, System X uses a high-performance switched communications fabric called Infiniband. Finally, it is cooled by a hybrid liquid-air cooling system. System X was assembled in only three weeks for a very affordable price when compared to other world-class supercomputers. The speaker will present the motivation for System X, describe its architecture and discuss the challenges faced in building, deploying and maintaining a large-scale supercomputer.

Srinidhi Varadarajan received his Ph.D. in Computer Science from the State University of New York, Stony Brook in 2000. He is the Director of the Terascale Computing Facility at Virginia Tech and an Assistant Professor in the Department of Computer Science. Dr. Varadarajan is the architect of System X, the third fastest supercomputer in the world, located at the Terascale Computing Facility. Dr. Varadarajan's research is focused on transparent fault tolerance for massively parallel supercomputers, scaleable network emulation,



Langley Research Center

iNASA


(**i**ntegrated **N**ASA **A**rchitecture
for **S**imulation and **A**nalysis)

- iNASA initiative will tackle NASA's ***hardest Science and Engineering Modeling and Simulation Challenges*** enabling unprecedented advances in multi-physics simulations with increased fidelity and resolution.
 - In concert with unique data created by science payloads, iNASA will create the capability to **model and understand** the processes that drive change in the Earth and in the Cosmos
 - Will allow, on a **day-to-day basis**, the capacity and capability to tackle simultaneously, at the limit of our engineering ability, problems that heretofore were untractable except on an **episodic/emergency basis**

HPC Across NASA Enterprises

Space Science

Solar Modeling




Exploration Systems

Crew Exploration Vehicle



Earth Science

Climate Modeling




Biological and Physical Research

Digital Astronaut




Space Flight

Emergency Response



Engineering and Safety

Shuttle Stewardship



Aeronautics

Vehicle Design



Each mission requires:

- Advanced models
- Model integration
- Efficient codes
- TF/PF Computation
- TB/PB data analysis, exploration, management
- Ensemble Analysis
- Science/engineering environments
- Remote resource access




Exploring **Scientific** Applications on Reconfigurable Hypercomputers



62K gates/FPGA


'02

Creativity & Innovation



6M gates/FPGA
=> 60M soon

'04





Langley Research Center

Computing **Faster** Without CPUs

GOAL: Evaluate FPGA-based Hypercomputers
for NASA Scientific Computations

TEAM: Drs. Olaf Storaasli & Jarek Sobieski, Principal Investigators
Dr. Robert Singleterry, Dave Rutishauser, Joe Rehder, Garry Qualls
Shaun Foley-*MIT*, William Fithian-*Harvard*, Siddhartha Krishnamurthy-*VT*
Cris Kania-*VT*, Patrick Butler-*VT*, Hoy Loper-*GS*, Neha Dandawate-*UVA*,
Kristin Barr-*JPMorgan*, Robert Lewis-*Morehouse*, Vincent Vance-*VT* +2 more


PARTNERS: Starbridge Systems, NSA, USAF, MSFC, AlphaStar



Langley-Starbridge Systems (SBS) Collaboration

- 1st SBS Hypercomputers @Langley (β -site: early H/W & S/W)
~ 1st Cray Supercomputer @LANL (8MB, 160MF, no O/S)
- **Langley:**
 - *application utilities & algorithms*
 - *identify bugs, recommend changes & utilities to SBS*
 - *coordinate with partners: NSA, USAF, MFSC, VPI*
SC89 (GigaFLOP Award) Rollwagen “NASA saved Cray \$2M”
- **Starbridge:**
 - *implement NASA suggestions*

==> Make FPGAs better tool for NASA applications



FPGA Computing

MetaLib

Traditional CPU

Sequential: 1 operation/cycle

Fixed gates & data types

Wasteful: 99% gates idle/cycle
yet all draw power

Software: Text

do i = 1, billion

c = a+b

end do

26 MFLOPS/250 MHz SGI

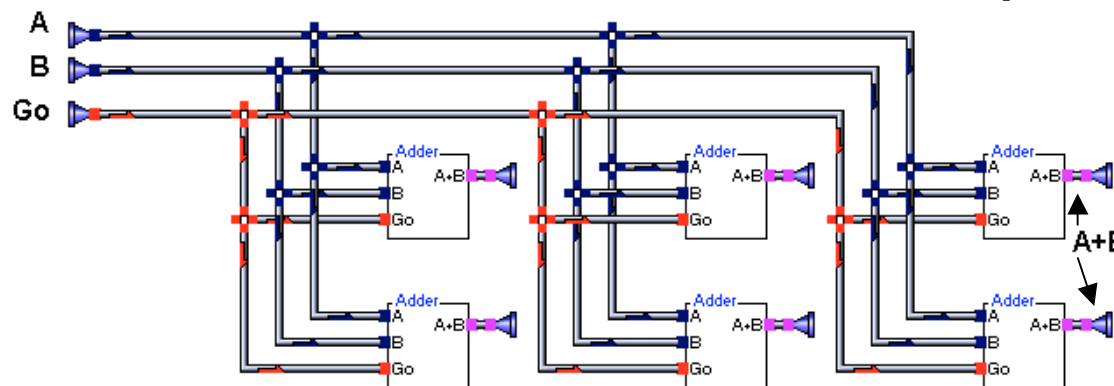
FPGA

Parallel: inherent

Dynamic gates & data types

Efficient: Optimizes gates to task

Gateware: VIVA Icons & Transports



47+ GFLOPS/64 MHz FPGA

329+ GFLOPS/10 FPGA board

- New Project
- MetaLib
- Basic Data Sets
- COM Data Sets
- Primitive Objects
 - Input
 - Output
 - \$Select
 - AND
 - DeRef
 - INVERT
 - OR
 - Ref
 - Release
 - Text
- Composite Objects
 - Control
 - Convert
 - DataInfo
 - ExposeCollect
 - Gates
 - I/O
 - Math
 - Memory
 - Mux
 - Registers
 - Shifting
 - TestSheets

I2ADL Editor Data Set Editor

Project System COM

Start


C:\Viva

Viva 2.0

2:58 PM

Algorithms Developed*

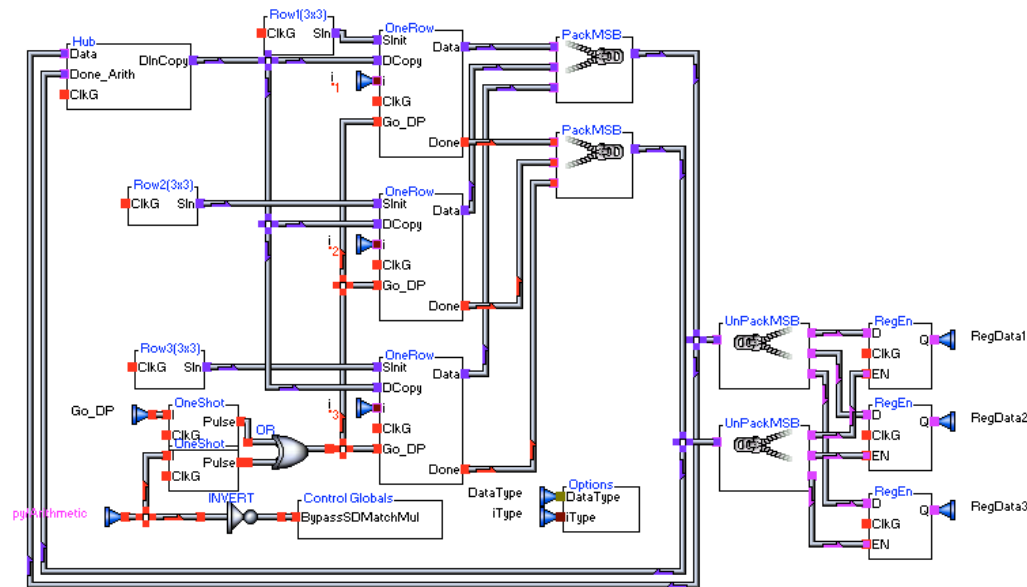
- **Matrix Algebra**: $\{V\}$, $[M]$, $\{V\}^T\{V\}$, $[M] \times [M]$, GCD, \dots
 - **$n!$** \Rightarrow Probability: Combinations/Permutations
 - **Cordic** \Rightarrow Transcendentals: \sin , \log , \exp , $\cosh \dots$
- $\partial y / \partial x$ & $\int f(x) dx \Rightarrow$ Runge-Kutta: CFD, Newmark Beta: CSM
- **Matrix Equation Solvers**: $[A]\{x\} = \{b\}$ - Gauss & Jacobi
 - **Dynamic Analysis**: $[M]\{\ddot{u}\} + [C]\{\dot{u}\} + [K]\{u\} + \text{NLT} = \{P(t)\}$
- **Analog Computing**: digital accuracy
 - **Nonlinear Analysis**: Analog simulation
avoids Non-Linear Term solution time



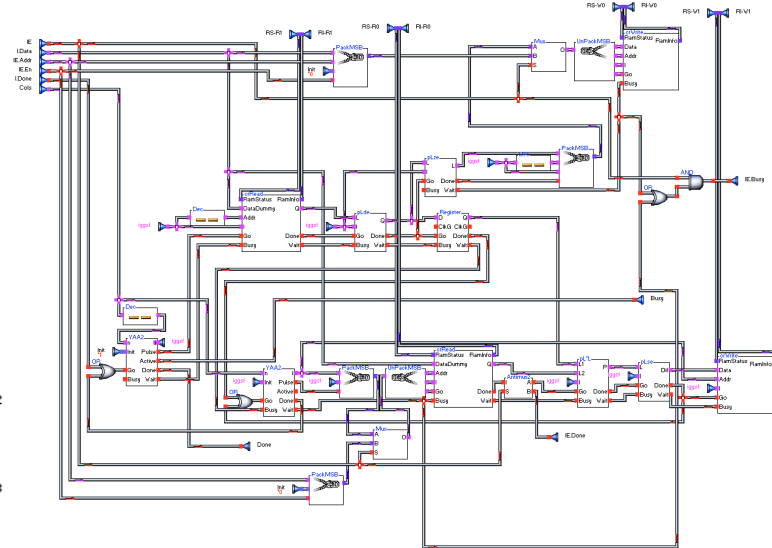
*Google: Publications Storaasli

Applications: VIVA Code

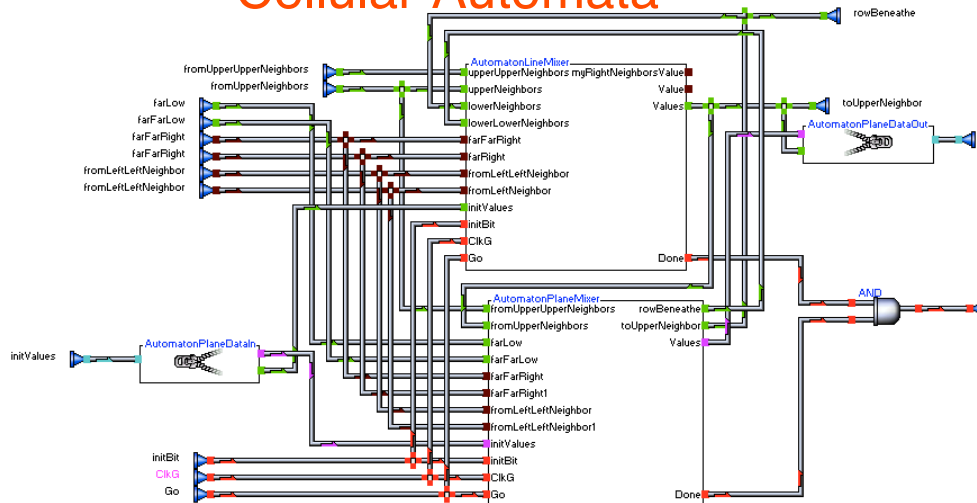
Jacobi Matrix Solver




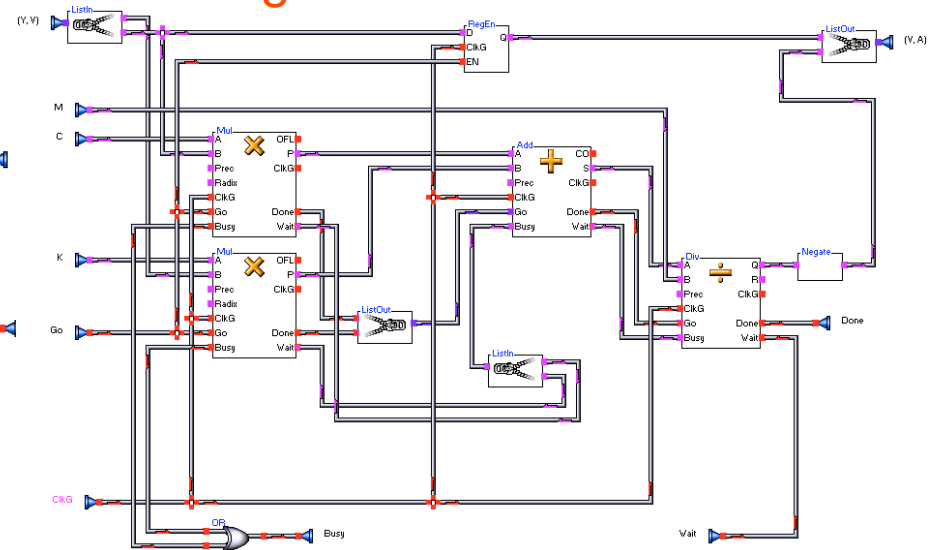
Gauss Matrix Solver



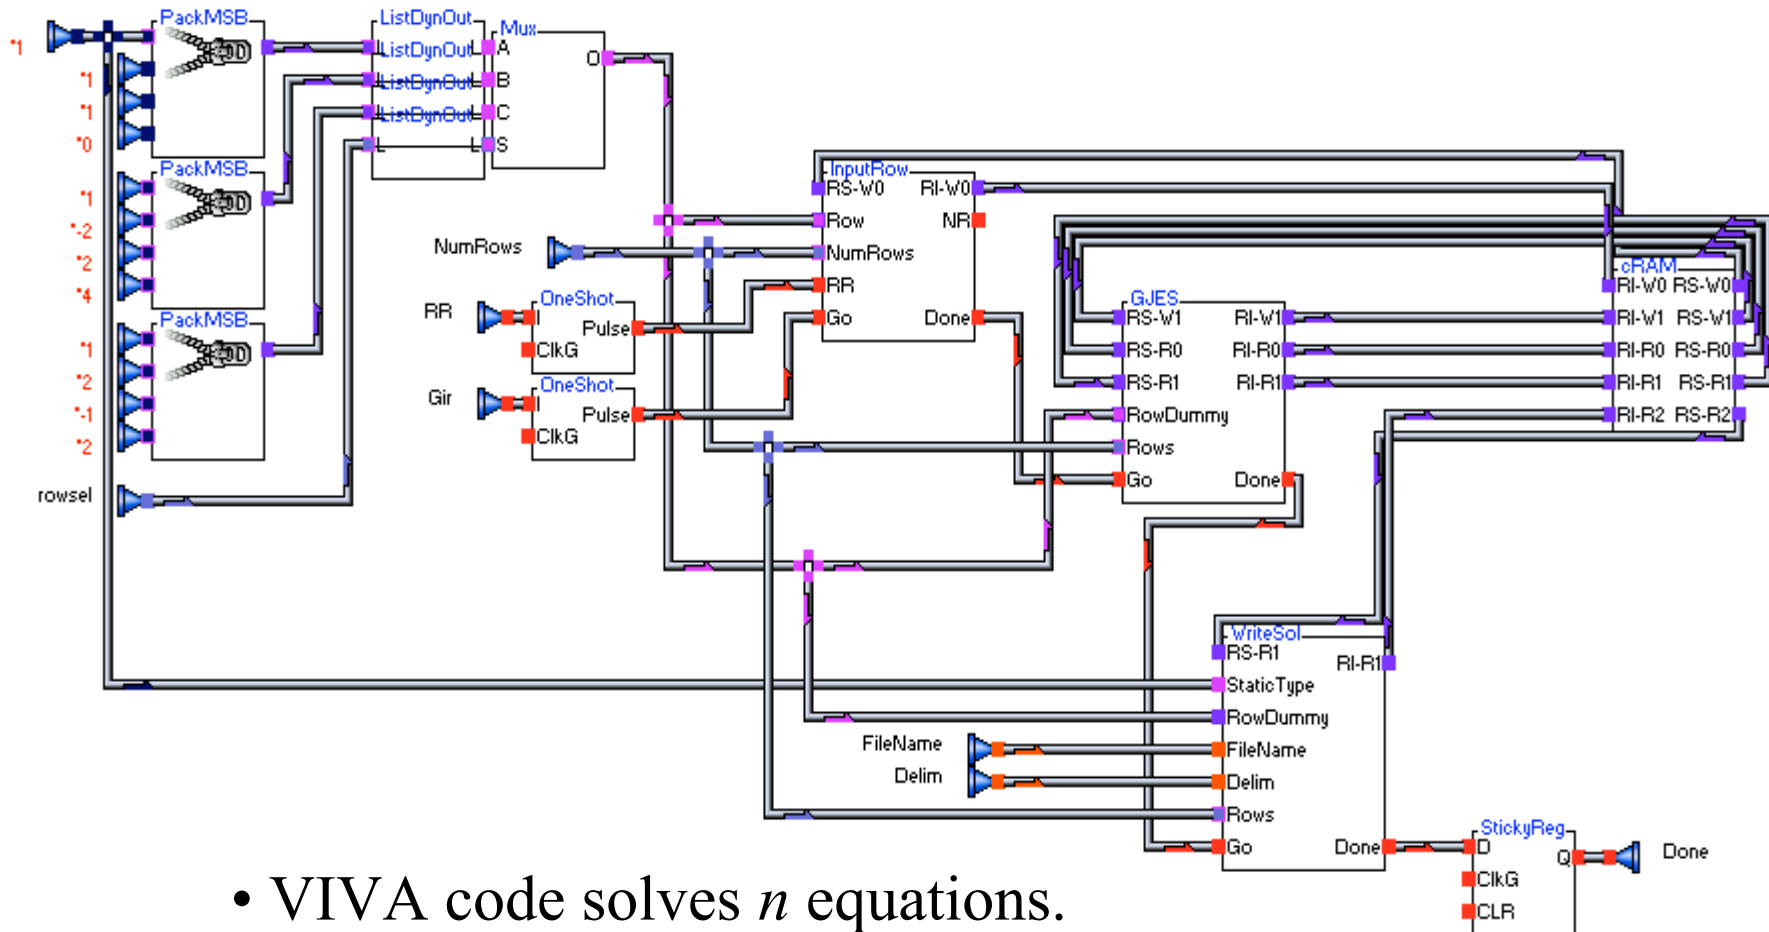
Cellular Automata



Runge-Kutta




Gauss-Jordan $\mathbf{A}x = \mathbf{B}$ Solver




- VIVA code solves n equations.

$$\begin{array}{rcl} \text{Ex: } x_0 + x_1 + x_2 & = & 0 \\ x_0 - 2x_1 + 2x_2 & = & 4 \\ x_0 + 2x_1 - x_2 & = & 2 \end{array} \Rightarrow \begin{array}{l} x_0 = 4 \\ x_1 = -2 \\ x_2 = -2 \end{array}$$

- Run on hypercomputer emulator, then FPGA



Spring-Mass Solver



Method: 4-stage Runge-Kutta

$$\frac{du}{dt} = f(u, t)$$

$$k_1 = hf(x_n, y_n)$$

$$k_2 = hf(x_n + \frac{1}{2}h, y_n + \frac{1}{2}k_1)$$

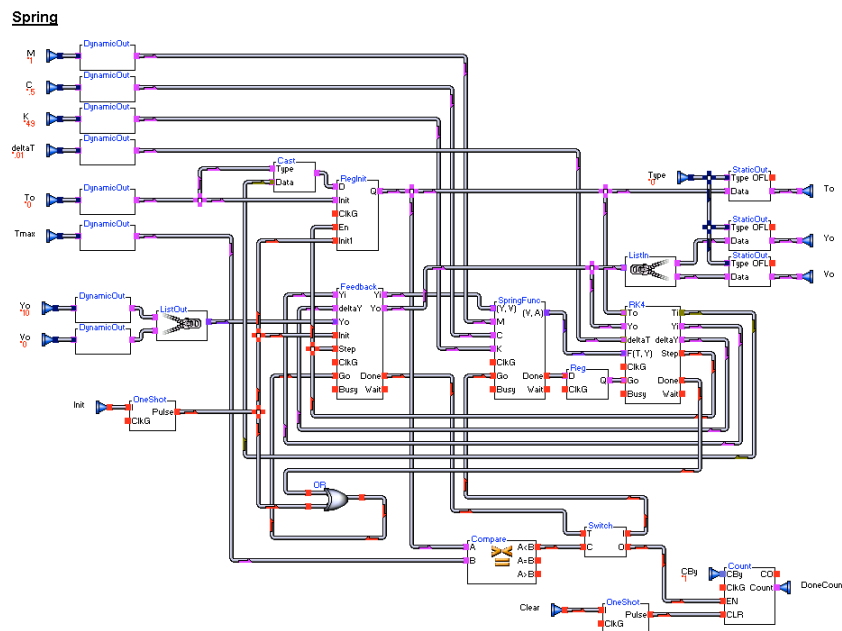
$$k_3 = hf(x_n + \frac{1}{2}h, y_n + \frac{1}{2}k_2)$$

$$k_4 = hf(x_n + h, y_n + k_3)$$


$$y(x_0) = y_0$$

$$x_{n+1} = x_n + h$$

$$y_{n+1} = y_n + \frac{1}{6}(k_1 + 2k_2 + 2k_3 + k_4)$$



Displacement vs. Time

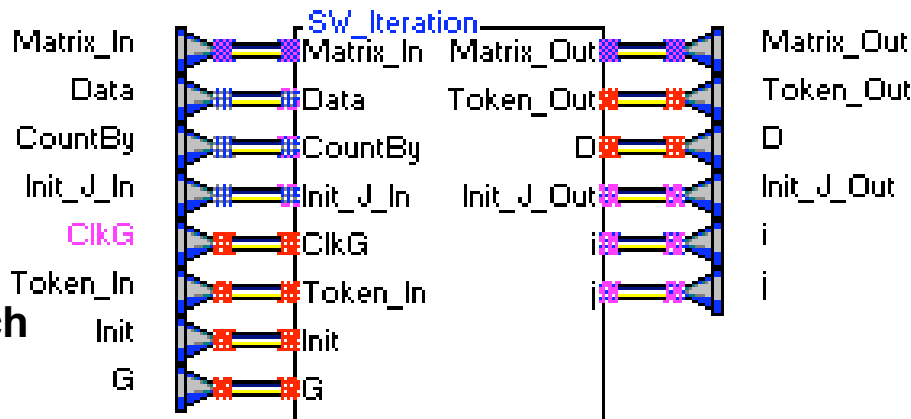


Smith Waterman Algorithm

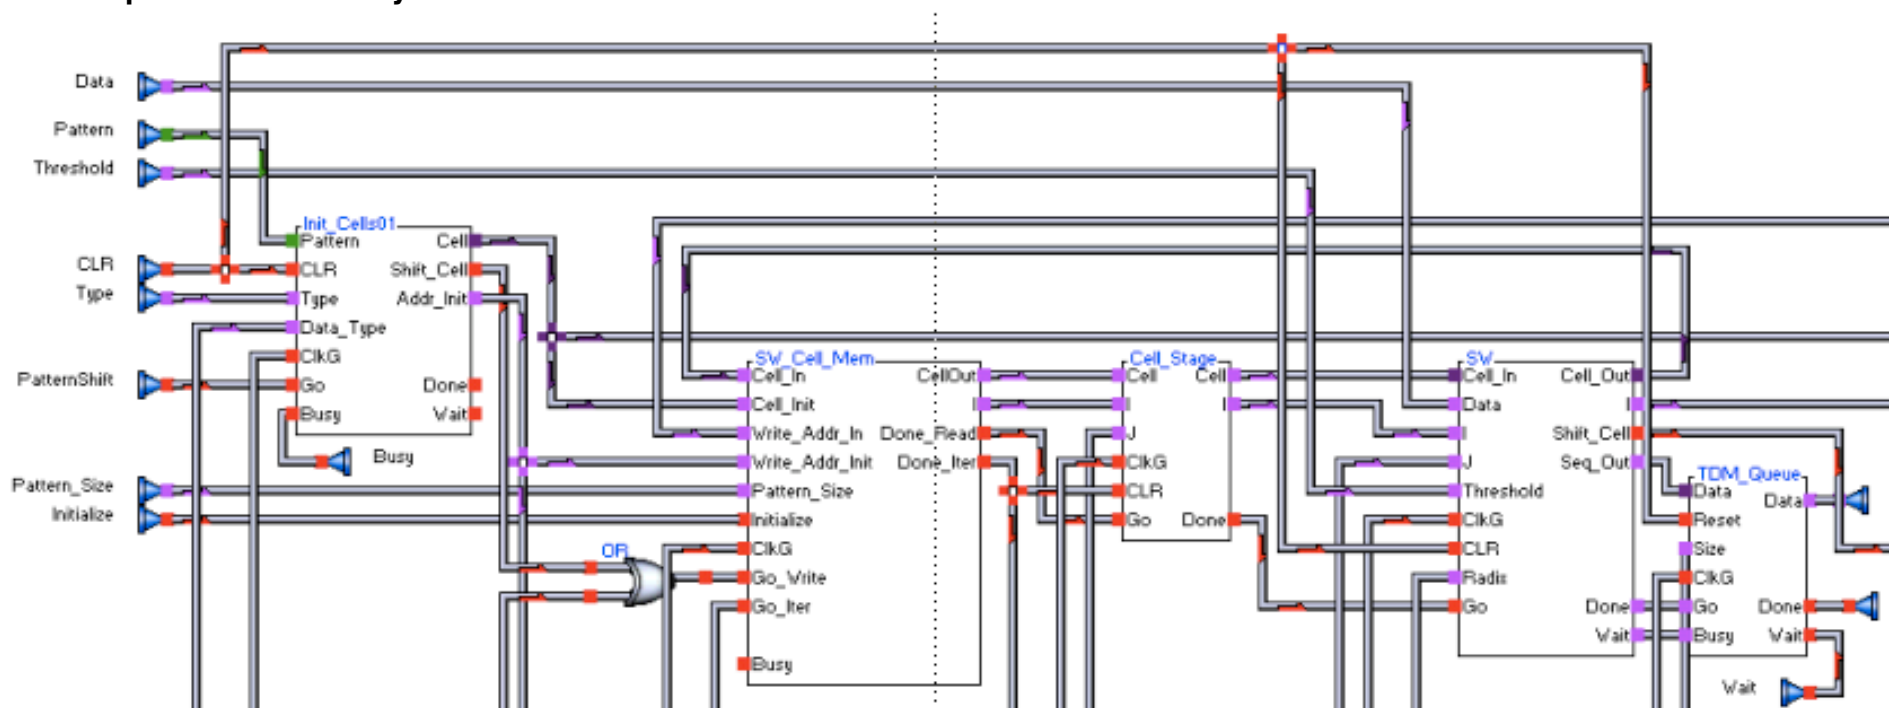
(NCI Compares 2 DNA sequences)

stream of cells*
 ref character
 # iterations, radix
 iteration index
 system clock
 pulse before cells
 pulse before search
 for each cell

* Keep full for efficiency



newly-computed cell
 for next iteration
 new cell
 for next iteration
 current row
 current column



FPGA SW Performance

7 FPGAs, on HC62, compute SW cells in parallel

(each FPGA computes 64 iterations/cycle)

⇒ 448 parallel SW_Iterations @ 25 Mhz*

⇒ $25\text{M} * 448 = \mathbf{11.2 \text{ billion SW steps/sec}}$


* *conservative: may increase clock via VIVA*

VIVA adds on 1 FPGA

Adds-parallel	16	32	128	256	512	640
% FPGA used	1	2	8	16	41	51
Billion Ops	4	8	34	77	154	192

⇒ **1000 adds/clock cycle**

(most CPUs limited to 1-3 ops/cycle)



Summary

NASA exploits HPC mix

- ✓ Production & Research Computing
- ✓ NASA-wide & Center Clusters
- ✓ Traditional & “far out” G5s & FPGAs

HPC future bright at NASA as capability/\$ increases

More=>  : Publications Storaasli